

REMARKS

Claims 1-20 were previously pending in this patent application.

Claims 1-20 stand rejected. Herein, Claims 1, 8, 12, and 17 have been amended. Accordingly, after this Amendment and Response, Claims 1-20 remain pending in this patent application. Further examination and reconsideration in view of the claims, remarks, and arguments set forth below is respectfully requested.

OATH/DECLARATION

It was stated in the Office Action (at page 2) that the oath or declaration was defective because no oath or declaration had apparently been provided.

Herein, a copy of the Declaration and Power of Attorney filed with present patent application on 02/20/2002, a copy of the postcard filed with present patent application on 02/20/2002, and a copy of the Patent Application Transmittal letter filed with present patent application on 02/20/2002, have been attached as proof that a signed declaration was filed with the present patent application on 02/20/2002. The Declaration and Power of Attorney is signed by the applicant and indicates that the present

patent application is attached to it on the filing date 02/20/2002. The postcard lists a Declaration/Power of Attorney as filed with the present patent application on 02/20/2002. Further, the Patent Application Transmittal letter indicates that a signed declaration and Power of Attorney is enclosed with the filed with the present patent application on 02/20/2002. Therefore, the declaration is not defective.

SPECIFICATION

The abstract of the disclosure is objected to because "is selected to replace" should be changed to "is selected to be replaced" and because "is configured to redirect" should be changed to "is subsequently configured to redirect". Herein, the abstract has been amended as suggested. Therefore, it is respectfully requested that the objection to the abstract of the disclosure be withdrawn.

Further, the disclosure is objected to because of several informalities. Herein, the disclosure has been amended to correct the informalities. Moreover, the original specification has been replaced with the attached Substitute Specification pursuant to 37 C.F.R. Sections 1.121(b)(3) and 1.125. The Substitute Specification includes no new matter. A Clean

Version Specification has also been attached pursuant to 37 C.F.R.

Sections 1.121(b)(3) and 1.125. The Substitute Specification and the Clean Version Specification now have line numbers to facilitate amendments to the text and citations to the text.

Also, the disclosure is objected to because striping across memory banks is described while use of "memory word" in pages 9 and 10 [or paragraphs 0026-0027] implies no striping across memory banks. The disclosure describes a "memory word" as being divided into the memory banks. [See Abstract and pages 3, 5, and 6]. Further, the "memory word" is described as including data bits and ECC bits. [See page 6]. The Office Action at page 2 states that this description implies the existence of multiple parity memory banks among the non-spare memory banks in a stripe. Also, the Office Action at page 2 states that the disclosure describes at page 6 a method that "builds on the RAID 3 concept", implying the existence of just one parity memory bank in a stripe.

Irrespective of whether the disclosure implies the existence of no parity memory banks, at least one parity memory bank, or the existence of multiple parity memory banks, the phrase "builds on the RAID 3 concept" of page 6 [or paragraph 0019] does not conflict because the disclosure clearly

states at page 6 that the new swapping method mimics RAID 3's ability to reconstruct data from a turned off memory bank by introducing a spare memory bank to be idle until the hot swapping is to be performed. That is, the new swapping method introduces a spare memory bank to be idle until the hot swapping is to be performed. However, the disclosure does not state or imply that new swapping method utilizes the one parity memory bank scheme of RAID 3 or of any level of RAID.

Continuing, the Office Action at page 3 states that the disclosure discloses no striping across memory banks and describes each bank as containing an entire memory word at page 6 and 7 [or paragraphs 0026-0027]. Herein, page 6 of the disclosure has been amended to specify that a "memory word" is divided into the memory banks 60 and 70 for storage. Further, the "memory word" may include a plurality of data bits and a plurality of ECC bits. Page 6 of the disclosure also has been amended to specify that memory bank 60 stores a first portion of the memory word while the memory bank 70 stores a second portion of the memory word. Any reference to "memory word" in "memory bank" (e.g., memory bank 60, memory bank 70, or spare memory bank 80) is intended to refer to either the first portion or the second portion of the "memory word" stored in the memory bank (e.g., memory bank 60, memory bank 70, or spare memory

bank 80). That is, the "content" of each memory bank (e.g., memory bank 60, memory bank 70, or spare memory bank 80) is a plurality of portions of different memory words, not a plurality of entire memory words.

In light of these arguments, it is respectfully requested that the objection to the disclosure be withdrawn.

35 U.S.C. Section 112, Second Paragraph Rejections

Claims 1-20 stand rejected under 35 U.S.C. Section 112, Second Paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. In particular, it was stated that Claims 1, 8, 12, and 17 include the recitation "a memory word is divided into said memory banks" but does not correspond to the detailed disclosure.

The disclosure describes a "memory word" as being divided into the memory banks. [See Abstract and pages 3, 5, and 6]. Further, the "memory word" is described as including data bits and ECC bits. [See page 6]. Herein, page 6 of the disclosure has been amended to specify that a "memory word" is divided into the memory banks 60 and 70 for storage.

Further, the "memory word" may include a plurality of data bits and a plurality of ECC bits. Page 6 of the disclosure also has been amended to specify that memory bank 60 stores a first portion of the memory word while the memory bank 70 stores a second portion of the memory word. Any reference to "memory word" in "memory bank" (e.g., memory bank 60, memory bank 70, or spare memory bank 80) is intended to refer to either the first portion or the second portion of the "memory word" stored in the memory bank (e.g., memory bank 60, memory bank 70, or spare memory bank 80). That is, the "content" of each memory bank (e.g., memory bank 60, memory bank 70, or spare memory bank 80) is a plurality of portions of different memory words, not a plurality of entire memory words.

Further, the disclosure at page 6 states that the "content" of the selected memory bank 70 is compared with the "content" of the spare memory bank 80 such that correctable errors are ignored. For example, if the ECC scheme corrects single bit errors, single bit errors existing between the copy of each memory word (e.g., second portion of memory words) stored in the selected memory bank 70 and the copy of each memory word (e.g., second portion of memory words) stored in the spare memory bank 80 are ignored. This assumes that the first portion of the memory word stored in the memory bank 60 does not have a bit error that would increase the

number of bit errors to two and cause an uncorrectable error situation. If the first portion of the memory word stored in the memory bank 60 does not have a bit error, single bit errors existing in the entire memory word (combination of the first and second portions) which has data bits and ECC bits can be corrected by the ECC scheme that corrects single bit errors. Hence, ignoring correctable errors does not guarantee that the switch in memory banks will be successful, but it will eliminate most failing cases (e.g., when the spare memory bank 80 is grossly defective).

Additionally, it was stated that Claim 4 lacks support in the disclosure. Herein, the disclosure at page 14 has been amended to specify that the process of hot swapping memory may further be repeated to replace every memory bank in turn, clearly supporting Claim 4.

In light of these arguments, it is respectfully requested that the rejection under 35 U.S.C. Section 112, Second Paragraph, against Claims 1-20 be withdrawn.

35 U.S.C. Section 103(a) Rejections

Claims 1-4, 7-9, 12-13, and 16-18 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Parks et al., U.S. Patent No. 6,598,174 (hereafter Parks) in view of Official Notice. These rejections are respectfully traversed.

Independent Claim 1 recites:

A method of hot swapping memory, comprising:

- a) providing a spare memory bank in a memory system, wherein said memory system includes a plurality of memory banks such that a memory word is divided into said memory banks;
- b) selecting one of said memory banks to replace;
- c) configuring said memory system to perform write operations associated with said selected memory bank to both said selected memory bank and said spare memory bank;
- d) ***performing atomic read and write operations such that content of said selected memory bank is copied directly to said spare memory bank***; and
- e) configuring said memory system to redirect operations to be performed on said selected memory bank to said spare memory bank such that said selected memory bank can be hot replaced.
(emphasis added)

It is respectfully asserted that Parks and the Official Notice do not disclose the present invention as recited in Independent Claim 1. In particular, the Office action (at page 5) cites Col. 19, lines 45-59, as disclosing "performing atomic read and write operations such that content of said selected memory ... is copied directly to said spare memory ...". On the contrary, Parks is directed to a method unlike the method recited in Independent Claim 1. In particular, Parks discloses a method/process that

facilitates protection of data in, and replacement of, storage devices about to fail before the failure happens. [Parks; Abstract]. In operation, if the condition indicated by the warning event is sufficient to invoke replacement of the storage unit, then the process determines whether a dedicated spare unit is available (box 301). [Parks; Figure 3; Col. 8, lines 40-52]. If a spare is available, then it is allocated (box 303). Id. The data on the storage unit causing the warning event is migrated to the allocated replacement unit (box 304). Id.

Further, Parks states that the data migration utilizes an intermediate device (501). [Parks; Figures 5A, 5B, and 6; Col. 9, lines 24-67]. The intermediate device (501) includes memory for use as buffers. Id. Contents of the source storage device (511) are transferred through the intermediate device (501) into the destination storage device (521). Id. Additionally, Parks states that data migration starts with allocation of a buffer in the intermediate device (box 1311) to support a block transfer. [Parks; Figures 10-13; Col. 17, line 58 through Col. 18, line 5]. A copy of a first block in the source storage device is moved to the buffer (box 1312 and box 1313). Id. Next, the block is moved from the buffer to the destination storage device (box 1314 and box 1315). Id. That is, Parks is directed to indirectly copying/moving the content of the source storage device to the destination

storage device. However, Parks fails to disclose performing atomic read and write operations such that content of the selected memory bank is copied directly to the spare memory bank without going through an intermediate storage/buffer device.

Unlike Parks and the Official Notice, Independent Claim 1 is directed to a method of hot swapping memory. The method includes performing atomic read and write operations such that content of the selected memory bank is copied directly to the spare memory bank. While Parks is directed to indirectly copying/moving the content of the source storage device to the destination storage device via an intermediate storage/buffer device, Independent Claim 1 is directed to directly copying/moving the content of the selected memory bank to the spare memory bank. Therefore, it is respectfully submitted that Independent Claim 1 is patentable over Parks and the Official Notice and is in condition for allowance.

Dependent Claims 2-4 and 7 are dependent on allowable Independent Claim 1, which is allowable over Parks and the Official Notice. Hence, it is respectfully submitted that Dependent Claims 2-4 and 7 are patentable over Parks and the Official Notice for the reasons discussed above.

With respect to Independent Claim 8, it is respectfully submitted that Independent Claim 8 recites similar limitations as in Independent Claim 1. In particular, Independent Claim 8 is directed to a circuit. The circuit comprises a repeater coupled to a plurality of memory banks such that a memory word is divided into the memory banks and coupled to a spare memory bank. The repeater directs write operations to be performed on a selected memory bank to both the selected memory bank and the spare memory bank. After atomic read and write operations are performed such that content of the selected memory bank is copied directly to the spare memory bank, the repeater redirects operations to be performed on the selected memory bank to the spare memory bank such that the selected memory bank can be hot replaced. Therefore, Independent Claim 8 is allowable over Parks and the Official Notice for reasons discussed in connection with Independent Claim 1.

Dependent Claim 9 is dependent on allowable Independent Claim 8, which is allowable over Parks and the Official Notice. Hence, it is respectfully submitted that Dependent Claim 9 is patentable over Parks and the Official Notice for the reasons discussed above.

With respect to Independent Claim 12, it is respectfully submitted that Independent Claim 12 recites similar limitations as in Independent Claim 1. In particular, Independent Claim 12 recites a memory system comprising a plurality of memory banks such that a memory word is divided into the memory banks; and a spare memory bank. Write operations associated with a selected memory bank are directed to both the selected memory bank and the spare memory bank, wherein atomic read and write operations are performed such that content of the selected memory bank is copied directly to the spare memory bank. Operations to be performed on the selected memory bank are redirected to the spare memory bank such that the selected memory bank can be hot replaced. Therefore, Independent Claim 12 is allowable over Parks and the Official Notice for reasons discussed in connection with Independent Claim 1.

Dependent Claims 13 and 16 are dependent on allowable Independent Claim 12, which is allowable over Parks and the Official Notice. Hence, it is respectfully submitted that Dependent Claims 13 and 16 are patentable over Parks and the Official Notice for the reasons discussed above.

With respect to Independent Claim 17, it is respectfully submitted that Independent Claim 17 recites similar limitations as in Independent Claim 1. In particular, Independent Claim 17 recites a computer system comprising a memory system including a plurality of memory banks such that a memory word is divided into the memory banks, a spare memory bank, and a repeater coupled to the memory banks and the spare memory bank. Write operations associated with a selected memory bank are directed to both the selected memory bank and the spare memory bank, wherein atomic read and write operations are performed such that content of the selected memory bank is copied directly to the spare memory bank. Operations to be performed on the selected memory bank are redirected to the spare memory bank such that the selected memory bank can be hot replaced. Therefore, Independent Claim 17 is allowable over Parks and the Official Notice for reasons discussed in connection with Independent Claim 1.

Dependent Claim 18 is dependent on allowable Independent Claim 17, which is allowable over Parks and the Official Notice. Hence, it is respectfully submitted that Dependent Claim 18 is patentable over Parks and the Official Notice for the reasons discussed above.

Claims 2, 10, 14, and 19 stand rejected under 35 U.S.C. 103(a) as being unpatentable over Parks et al., U.S. Patent No. 6,598,174 (hereafter Parks) in view of Official Notice, and further in view of Ohizumi, U.S. Patent No. 5,357,509 (hereafter Ohizumi). These rejections are respectfully traversed.

Dependent Claim 2, Dependent Claim 10, Dependent Claim 14, and Dependent Claim 19 are dependent on allowable Independent Claims 1, 8, 12, and 17 respectively, which are allowable over Parks and the Official Notice. Moreover, Ohizumi does not disclose performing atomic read and write operations such that content of the selected memory bank is copied directly to the spare memory bank, as recited in Claims 1, 8, 12, and 17. On the contrary, Ohizumi is directed to generating restored data from data stored in remaining functioning disks, writing the restored data to the spare disk storage, and copying the restored data to a new disk storage which has replaced the faulty disk. Hence, it is respectfully submitted that Dependent Claims 2, 10, 14, and 19 are patentable over Parks, the Official Notice, and Ohizumi for the reasons discussed above.

ALLOWABLE SUBJECT MATTER

Claims 3-6, Claim 11, Claim 15, and Claim 20 would be allowable if rewritten to overcome the rejections under 35 U.S.C. Section 112, Second Paragraph, and to include all of the limitations of the base claim and any intervening claims.

Claims 3-6, Claim 11, Claim 15, and Claim 20 are dependent on allowable Independent Claims 1, 8, 12, and 17 respectively, which are allowable. Hence, it is respectfully submitted that Dependent Claims 3-6, 11, 15, and 20 are patentable.

CONCLUSION

It is respectfully submitted that the above arguments and remarks overcome all rejections. For at least the above-presented reasons, it is respectfully submitted that all remaining claims (Claims 1-20) are now in condition for allowance.


The Examiner is urged to contact Applicant's undersigned representative if the Examiner believes such action would expedite resolution of the present Application.

Please charge any additional fees or apply any credits to our PTO deposit account number: 23-0085.

Respectfully submitted,

Wagner, Murabito & Hao, LLP

Dated: 11/17/04



John P. Wagner
Registration No. 35,398

Two North Market Street, Third Floor
San Jose, CA 95113
(408) 938-9060

Attachments



HOT SWAPPING MEMORY METHOD AND SYSTEM

BACKGROUND OF THE INVENTION

5 FIELD OF THE INVENTION

The present invention generally relates to computer systems. More particularly, the present invention relates to memory systems.

RELATED ART

10 To improve reliability, availability, and serviceability, a variety of techniques have evolved to facilitate hot swapping memory in computer system such as personal computers and servers. This allows the memory defect (or failing memory) to be healed (or replaced) without taking the computer system down. Moreover, substantial error correction capability
15 has been integrated into servers, allowing them to run with a faulty memory module without crashing.

Traditionally, hot swapping memory has been accomplished by mirroring. That is, a second copy of the memory content is provided in the
20 main memory system. For every memory bank in the main memory system, there exists a mirror memory bank having the same content. Every write operation to the main memory writes two copies: one copy to the memory bank and one copy to the mirror memory bank. Each read comes from a single copy of the main memory system.

25

Many implementations read just one copy at a time--if the copy being read has an uncorrectable error (through whatever error correction code (ECC) scheme that is being used), the computer system will report an uncorrectable error and crash even though there probably is a correct copy
30 of the read in the unread memory copy. This is an implementation

SUBSTITUTE SPECIFICATION

100200334-1

optimization. The number of ECC corrections can be used as a trigger to switch which copy from main memory is being read at any particular time.

A hot swapping operation is accomplished by suspending all
5 accesses to a memory bank (mirror or non-mirror), and then turning that
memory bank off. Certain maintenance operations are done in order to
make sure that both the memory bank and the mirror memory bank are
consistent, especially around hot swap operations. This is strongly
analogous to RAID 1 (redundant array of independent disks). It is easy to
10 implement, but quite expensive since two full copies of the contents of the
main memory are needed.

Another approach to hot swapping memory is based on RAID 3. In
this approach, the main memory system has one copy plus some extra
15 information to help recover if a small portion of the main memory fails.
Typically, this is accomplished by dividing the main memory system into
several memory banks, striping the data across the memory banks, and
adding one extra memory bank that stores the parity (or some other
function) of the data stored in the other memory banks. In this way, if the
20 failing memory bank is known, the failing memory bank can be
reconstructed from the remaining memory banks and the extra memory
bank storing the parity information. This has the advantage that less
memory capacity is needed than the mirroring approach, but at the cost of a
more complex algorithm (e.g., to calculate parity) for managing the main
25 memory system.

SUBSTITUTE SPECIFICATION

100200334-1

SUMMARY OF THE INVENTION

A method of hot swapping memory is described. A memory system includes a plurality of memory banks such that a memory word is divided into the memory banks. The memory system is provided a spare memory bank. One of the memory banks is selected to ~~replace~~ be replaced. The memory system is configured to perform write operations associated with the selected memory bank to both the selected memory bank and the spare memory bank. Moreover, atomic read and write operations are performed such that the content of the selected memory bank is copied to the spare memory bank. Furthermore, the memory system is subsequently configured to redirect operations to be performed on the selected memory bank to the spare memory bank such that the selected memory bank can be hot replaced.

SUBSTITUTE SPECIFICATION

100200334-1

BRIEF DESCRIPTION OF THE DRAWINGS

5 The accompanying drawings, which are incorporated in and form a part of this specification, illustrate embodiments of the invention and, together with the description, serve to explain the principles of the present invention.

Figure 1 illustrates a block diagram of a computer system in accordance with an embodiment of the present invention.

10 Figure 2 illustrates a flow chart showing a method of hot swapping memory in accordance with an embodiment of the present invention.

Figures 3A-3F illustrate memory data flow in accordance with an embodiment of the present invention.

15

The drawings referred to in this description should not be understood as being drawn to scale except if specifically noted.

100200334-1

DETAILED DESCRIPTION OF THE INVENTION

Reference will now be made in detail to the preferred embodiments of the present invention, examples of which are illustrated in the accompanying drawings. While the invention will be described in conjunction with the
5 preferred embodiments, it will be understood that they are not intended to limit the invention to these embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents, which may be included within the spirit and scope of the invention as defined by the appended claims. Furthermore, in the following detailed description of the
10 present invention, numerous specific details are set forth in order to provide a thorough understanding of the present invention.

HOT SWAPPING MEMORY

Figure 1 illustrates a block diagram of a computer system 100 in
15 accordance with an embodiment of the present invention. As illustrated in Figure 1, the computer system 100 includes a chipset 40, one or more processors 20, one or more input/output data ports 30, and a memory system 50. In an embodiment, the memory system 50 is a main memory system 50. The chipset 40 interfaces the processor(s) 20 with the
20 input/output data port(s) 30 and the main memory system 50. It should be understood that the computer system 100 can have other configurations.

In an embodiment, the main memory system 50 includes one or more repeaters 10 coupled to a plurality of memory banks 60 and 70 and coupled
25 to a spare memory bank 80, whereas a memory word is divided into the memory banks 60 and 70 for storage. The repeater 10 can have separate circuit modules for the spare memory bank 80 and each memory bank 60 and 70 to facilitate write operations and read operations to the spare memory bank 80 and the memory banks 60 and 70. The spare memory
30 bank 80 and each memory bank 60 and 70 include one or more memory

100200334-1

modules 5, whereas each memory module 5 includes one or more memory chips. In an embodiment, the memory modules 5 are dual in-line memory modules (DIMMs). It should be understood that the memory module 5 can be any other type of memory module. It should be understood that the main
5 memory system 50 can have more than two memory banks. Furthermore, it should be understood that the spare memory bank 80 and each memory bank 60 and 70 can have less than or more than four memory modules 5.

In an embodiment, the main memory system 50 implements any one
10 of a variety of error correction code (ECC) schemes. In the main memory system 50 implementing an ECC scheme, a memory word includes a plurality of data bits and a plurality of ECC bits. Moreover, each type of ECC scheme has a different error correction capability. For example, some ECC schemes provide for automatic correction when a single bit is in error
15 and provide for detection of two bits in error. Other ECC schemes provide multiple-bit correction. In particular, a chipkill ECC scheme enables the main memory system 50 to withstand a multi-bit failure within a memory chip of any one of the memory modules 5.

20 The computer system 100 supports a new hot swapping memory method, whereas the terms "hot swapping memory" refer to the capability to pull out or plug-in memory components (e.g., any of the memory banks 60 and 70 and the spare memory bank 80 of the computer system 100) while the computer system 100 is powered and still operating. The new hot
25 swapping memory method builds on the RAID 3 concept by combining RAID 3 concepts with ECC schemes associated with main memory. Thus, the new hot swapping memory method relies on the ECC scheme for data accuracy but mimics RAID 3's ability to reconstruct data from a turned off memory bank by introducing a spare memory bank to be idle until the hot
30 swapping is to be performed. Moreover, the new hot swapping memory

SUBSTITUTE SPECIFICATION

100200334-1

method is less costly than the mirroring approach discussed above since one spare memory bank is needed rather than a mirror memory bank for each memory bank of the main memory system 50. Also, the new hot swapping memory method avoids the complex algorithm associated with RAID 3 on every write to memory (and reads while a memory bank is failed/turned off).

In particular, the new hot swapping memory method is accomplished with minimal support from the computer system's 100 hardware and without complicated or time consuming operations that substantially interfere with the performance of the computer system 100 during normal operation. The hot swapping memory can be implemented with hardware within the repeater 10. Alternatively, the hot swapping memory can be implemented with hardware within the chipset 40 or within any other location in the computer system 100.

In practice, the repeater 10 increases the memory capacity of the main memory system 50 and may have multiplexing capability. The repeater 10 may be implemented as a bit-sliced repeater that receives some bits from every memory bank. To support the new hot swapping memory method, the repeater 10 is configured to direct write operations for the main memory system 50 to a memory bank (e.g., memory bank A 60 or memory bank B 70), to a spare memory bank 80, or to both a memory bank (e.g., memory bank A 60 or memory bank B 70) and a spare memory bank 80. Moreover, the repeater 10 is configured to direct read operations for the main memory system 50 to a memory bank (e.g., memory bank A 60 or memory bank B 70) or to a spare memory bank 80.

Figure 2 illustrates a flow chart showing a method 200 of hot swapping memory in accordance with an embodiment of the present

SUBSTITUTE SPECIFICATION

100200334-1

invention. Reference is made to Figures 1 and 3A-3F. Initially, the memory bank A 60, the memory bank B 70, and the spare memory bank 80 of the computer system 100 are in the operational states illustrated in Figure 3A. As depicted in Figure 3A, the memory bank A 60 and the memory bank B 70 are on-line, in ~~used~~ use, and populated with data received via the repeater 10. The arrow 310 indicates that read operations and write operations are being performed on the memory bank A 60 via the repeater 10. The arrow 320 indicates that read operations and write operations are being performed on the memory bank B 70 via the repeater 10. Moreover, the spare memory bank 80 is off-line via an isolation switch 330, is not being used for read operations or write operations, and is not populated with data. In fact, the spare memory bank 80 can be powered down to save power.

At Block 210 of Figure 2, one of the memory banks (e.g., memory bank A 60 or memory bank B 70) is selected to be replaced. The selection can be made based on any number of factors. For example, the selected memory bank may need to be upgraded, repaired, maintained, expanded, etc. Additionally, by monitoring correctable memory errors during the computer system's 100 normal operation, the selected memory bank may have accumulated a number of correctable memory errors that has exceeded a particular threshold. As depicted in Figure 3B, the memory bank B 70 has been selected to be replaced, whereas the arrow 340 indicates the selected memory bank.

Furthermore, at Block 220 of Figure 2, the main memory system 50 is configured to perform write operations associated with the selected memory bank 70 to both the selected memory bank 70 and the spare memory bank 80. In an embodiment, the repeater 10 directs write operations associated with the selected memory bank 70 to both the selected memory bank 70 and the spare memory bank 80. As depicted in Figure 3C, the spare

SUBSTITUTE SPECIFICATION

100200334-1

memory bank 80 is on-line but is not being used. Moreover, the arrow 350 indicates that read operations associated with the selected memory bank 70 are being performed on the selected memory bank 70. However, the arrow 360 indicates that write operations associated with the selected memory
5 bank 70 are being performed on the selected memory bank 70 and the spare memory bank 80.

At Block 230 of Figure 2, atomic read and write operations are performed such that the content of the selected memory bank 70 is copied
10 to the spare memory bank 80. Normal memory accesses to the memory banks 60 and 70 continue during these atomic read and write operations. Any reduction in performance of the computer system 100 is dependent on the period of time in which these atomic read and write operations are performed. If these atomic read and write operations are performed in a
15 short period of time, there may be a reduction in the performance of the computer system 100. If these atomic read and write operations are performed in a longer period of time, there may be just a minimal reduction in the performance of the computer system 100. In Figure 3C, the arrow 370 indicates that atomic read and write operations are being performed.
20 The chipset 40, low level software, the repeater 10, or any other component such as a memory controller can be configured to scrub (i.e., perform atomic read and write operations) the selected memory bank 70 into the spare memory bank 80. For example, in an atomic operation, the memory controller reads the memory word in the memory banks 60 and 70, and
25 writes the memory word back into memory banks 60 and 70 and the spare memory bank 80. This is a common feature of memory controllers, and is intended to remove correctable soft errors from the main memory system
50.

SUBSTITUTE SPECIFICATION

100200334-1

At Block 240 of Figure 2, the content of the selected memory bank 70 is compared with the content of the spare memory bank 80 such that correctable errors are ignored. In an embodiment, the repeater 10 includes a comparator 390 (Figure 3D). As described above, a memory word is
5 divided into the memory banks 60 and 70 for storage. Further, the memory
word may include a plurality of data bits and a plurality of ECC bits. Thus,
memory bank 60 stores a first portion of the memory word while the memory
bank 70 stores a second portion of the memory word. Any reference to
10 "memory word" in "memory bank" (e.g., memory bank 60, memory bank 70,
or spare memory bank 80) is intended to refer to either the first portion or
the second portion of the memory word stored in the memory bank (e.g.,
memory bank 60, memory bank 70, or spare memory bank 80).

~~In particular~~ Continuing, a memory word of the selected memory
15 bank 70 and a memory word of the spare memory bank 80 are read and
compared until the entire content of the selected memory bank 70 is
compared with the entire content of the spare memory bank 80.

There are several types of correctable errors, whereas correctable
20 errors are bit errors that can be corrected by the ECC scheme implemented
by the main memory system 50 (Figure 1). Each type of correctable errors
error situation is dependent on the type of ECC scheme implemented by the
main memory system 50 (Figure 1). For example, if the ECC scheme
corrects single bit errors, the comparator 390 (Figure 3D) will ignore single
25 bit errors existing between the copy of the memory word stored in the
selected memory bank 70 and the copy of the memory word stored in the
spare memory bank 80. Thus, the spare memory bank 80 may not be error
free, but it will be good enough to work. More importantly, any errors
present in the selected memory bank 70 will not prevent the switch in
30 memory banks (i.e., from the selected memory bank 70 to the spare

100200334-1

memory bank 80) from occurring to facilitate hot replacing the selected memory bank 70. Similarly, if the ECC scheme is a chipkill ECC scheme or multibit ECC scheme, the comparator 390 (Figure 3D) will ignore bit errors existing in particular bit sets between the copy of the memory word stored in the selected memory bank 70 and the copy of the memory word stored in the spare memory bank 80. In Figure 3D, the arrow 380 indicates that a memory word is read from the selected memory bank 70 and sent to the comparator 390. Moreover, the arrow 385 indicates that a memory word is read from the spare memory bank 80 and sent to the comparator 390.

Continuing at Block 250 of Figure 2, it is determined whether the comparator 390 (Figure 3D) detected any uncorrectable errors. If the comparator 390 detected any uncorrectable errors, the method proceeds to Block 260. At Block 260, it is determined that the selected memory bank 70 cannot be hot replaced because the spare memory bank 80 is defective. Thus, the spare memory bank 80 must first be replaced. Then, the method 200 of Figure 2 can be restarted. As depicted in Figure 3D, the comparator 390 includes an indicator 395 for indicating the detection of uncorrectable errors. This is not enough to guarantee that the switch in memory banks (i.e., from the selected memory bank 70 to the spare memory bank 80 to facilitate hot replacing the selected memory bank 70) will be successful, but it will eliminate most of the failing cases (i.e., when the spare memory bank 80 is grossly defective).

Alternatively, the comparison operation (e.g., Blocks 240-260) can be omitted. However, performing the comparison operation (e.g., Blocks 240-260) increases the reliability of the switch in memory banks (i.e., from the selected memory bank 70 to the spare memory bank 80) to facilitate hot replacing the selected memory bank 70.

SUBSTITUTE SPECIFICATION

100200334-1

Otherwise, at Block 270 of Figure 2, if the comparison operation is successful, the main memory system 50 is configured to perform read and write operations associated with the selected memory bank 70 on the spare memory bank 80 rather than the selected memory bank 70. In an

5 embodiment, the repeater 10 redirects operations to be performed on the selected memory bank 70 to the spare memory bank 80 such that the selected memory bank 70 can be hot replaced. As depicted in Figure 3E, the selected memory bank 70 is on-line but is no longer being used. Thus, the selected memory bank 70 can be placed in an off-line state. Moreover,
10 the spare memory bank 80 is on-line and is being used for read operations and write operations. The arrow 400 indicates that the operations (read operations or write operations) to be performed on the selected memory bank 70 are being performed on the spare memory bank 80.

15 At Block 280 of Figure 2, the selected memory bank 70 is isolated and replaced without powering down the computer system 100. As depicted in Figure 3F, the selected memory bank 70 is off-line via the isolation switch 410 and is not being used. During the new hot swapping memory method of Figure 2, there was no need to turn off the computer system 100 or to limit
20 normal accesses to the main memory system 50 (Figure 1). The new hot swapping memory method of Figure 2 is dependent on the ECC scheme for error detection and correction but allows data to be copied between memory banks while normal memory accesses are occurring. In the new hot swapping memory method of Figure 2, the selected memory bank 70 is the
25 source memory bank while the spare memory bank 80 is the target memory bank.

In a dedicated spare memory bank embodiment, the selected memory bank 70 (or memory bank B 70) has to be replaced with a
30 functional memory bank and the content of the spare memory bank 80 has

100200334-1

to be copied to the functional memory bank using the new hot swapping memory method of Figure 2, before the memory bank A 60 can be selected to be replaced using the new hot swapping memory method of Figure 2.

5 For example, the selected memory bank 70 (which now is off-line and is not in use) is replaced with a functional memory bank. The spare memory bank 80 is selected such that the spare memory bank 80 is the source memory bank while the functional memory bank is the target memory bank. Then, the main memory system 50 is configured to perform write operations
10 associated with the spare memory bank 80 to both the spare memory bank 80 and the functional memory bank. Moreover, atomic read and write operations are performed such that content of the spare memory bank 80 is copied to the functional memory bank. Furthermore, the content of the spare memory bank 80 is compared with the content of the functional
15 memory bank such that correctable errors are ignored. Alternatively, the comparison operation can be omitted. However, performing the comparison operation increases the reliability of the switch in memory banks (i.e., from the spare memory bank 80 to the functional memory bank). If the comparison operation is successful, the main memory system 50 is
20 configured to redirect operations to be performed on the spare memory bank 80 to the functional memory bank. Thus, the spare memory bank (which now is off-line and is not in use) 80 can be used in the new hot swapping memory method of Figure 2 to hot replace any of the memory banks (e.g., memory bank A 60 or memory bank B 70).

25

 In a non-dedicated spare memory bank embodiment, any unused memory bank of the memory banks can be used in place of the spare memory bank 80. Thus, it is not necessary to copy the content of the spare memory bank 80 to another memory bank using the new hot swapping
30 memory method of Figure 2, before another memory bank can be selected

SUBSTITUTE SPECIFICATION

100200334-1

to be replaced using the new hot swapping memory method of Figure 2. In case that the spare memory bank 80 is being used and is populated with data, any memory bank (e.g., memory bank A 60 or memory bank B 70) that is not populated with data and is not being used can be utilized in place of the spare memory bank for the new hot swapping memory method of Figure 2.

For instance, the selected memory bank 70 (which now is off-line and is not in use) is replaced with a functional memory bank. Then, a particular memory bank from the memory bank A 60 and the spare memory bank 80 is selected to be hot replaced such that the particular memory bank is the source memory bank while the functional memory bank is the target memory bank. The main memory system 50 is configured to perform write operations associated with the particular memory bank to both the particular memory bank and the functional memory bank. Then, atomic read and write operations are performed such that the content of the particular memory bank is copied to the functional memory bank. The content of the particular memory bank is compared with the content of the functional memory bank such that correctable errors are ignored. Alternatively, the comparison operation can be omitted. However, performing the comparison operation increases the reliability of the switch in memory banks to facilitate hot replacing a memory bank. If the comparison operation is successful, the main memory system 50 is configured to redirect operations to be performed on the particular memory bank to the functional memory bank.

In order to facilitate further hot replacing of other memory banks, the particular memory bank (which now is off-line and is not in use) is replaced with a second functional memory bank. In a similar manner as described above, a second particular memory bank from the memory banks and the spare memory bank 80 is selected to be replaced. Moreover, the new hot

SUBSTITUTE SPECIFICATION

100200334-1

swapping memory method of Figure 2 can be performed using the second functional memory bank as the target memory bank and using the second particular memory bank as the source memory bank. This process of hot swapping memory may further be repeated to replace every memory bank in turn.

The foregoing descriptions of specific embodiments of the present invention have been presented for purposes of illustration and description. They are not intended to be exhaustive or to limit the invention to the precise forms disclosed, and many modifications and variations are possible in light of the above teaching. The embodiments were chosen and described in order to best explain the principles of the invention and its practical application, to thereby enable others skilled in the art to best utilize the invention and various embodiments with various modifications as are suited to the particular use contemplated. It is intended that the scope of the invention be defined by the Claims appended hereto and their equivalents.